

Articles

Elements of A.I.

Artificial intelligence / Online course

In the next 6 weeks I will give here in individual chapters either my answers or a summary. This should help me personally to sort the topics of the course: Elements of A.I. [1] of the University of Helsinki [2] better, to learn and to save in mind. The exam question or my answer is not meant to be simply copied, but to inspire people outside the course to think. The difficult question of definition can only be presented in the light of the individual points proposed in the task. Developing a new definition from this requires more time for current developments and results. So far there has been no very precise definition, either scientifically or socially.

The definition “cool things that computers can't do” is already largely correct. A self driving car [3] with an A.I. [4] method is more likely to plan current traffic through computer vision [5]. A computer could only follow a fixed route and no longer reliably follow it if something unforeseen should happen. This also applies to other autonomous systems, such as delivery robots, drones or ships [6]. In the area of image and video processing, AI methods also comply with the definition. A Photoshop filter is a fixed filter that cannot be changed, but image style transfer [7] can combine different presets with one another and thus edit a wider range of tasks. They should therefore rather be described as autonomous and adaptive systems.

I find the definition *machines imitating intelligent human behavior* rather weak, because a robot (e.g. in the form of a human) can imitate human behavior. Alan Turing [8] considered this question with his Turing Test [9], which was attacked by John Searle [10] with his Chinese Room counterargument [11]. It was also about the question of which intelligence should be defined for artificial intelligence. A customer chatbot can certainly not be described as intelligent, because he only reacts to questions and tries to answer them as well as possible. One does not know how a human being can react intelligently and intelligence cannot be measured either, just like a single dimension e.g. the temperature in the bathtub or being able to solve a problem tells us nothing about the ability of solve another, different problem.

Preface:

If you are learning, look for a very quiet place where you can learn. I made two small mistakes, which were not made because I did not understand that, but because my concentration was distracted by something. In this chapter you have to sit down with pen and paper to go through everything step by step and you don't need a distraction. A little hint. Some puzzles have more than one solution and sometimes the obvious is not always the best.

The chapter is about *search in practice* e.g. for self-propelled cars and games. You stick to the topic of games and why they can be connected to AI. The first example is about the Chicken Crossing River Puzzle [12] and how the puzzle is constructed, what the starting conditions are and where the solution must lead. The three points of state space, transitions and costs are discussed in more detail. The test is to solve the problem and after going through it a few times on paper, it's not really difficult. Then comes the [Tower of Hanoi](#), which has to be searched through and worked through as an examination. This puzzle is also just an introduction, but one learns a lot about the basis.

In the second chapter a short excursion into the Second World War is made to visit Alan Turing again. John McCarty [13] is also discussed thematically and why he has to do with artificial intelligence. There is no exam in this chapter, but there is some background knowledge. The structure of the university is very well done, because the brain needs a short break after the two puzzles, before it goes to the last big subchapter.

The last subchapter is about Min and Max [14] playing the tic-tac-toe [15]. The tempo is tightened and you not only have to record a lot of information, but you also have to understand it. It is about game trees, minimizing and maximizing values, optimal move finding, problems of massive game trees, limitations of plain search and good heuristics. The check is quite overwhelming at first, but if you work through all subtasks in peace (and don't make any mistakes at the beginning) this is no problem. Overall, the chapter was exciting and the tasks stimulate further thinking. I also play with the idea to build the first two puzzles as games from real materials. Then you can play this with your friends, leave the tasks to them and explain again after the game. Because with me (personally) the best way to store information is to explain it in detail once again.

This chapter deals with artificial intelligence and how it can be applied in reality. In the subchapter odds and probability I made only slow progress and only scored 5/6 in the test. I have a hunch that I can't appreciate it very much (or I have to take a closer look at the concept). I had to work through The Bayers Rule [16] a little and try hard and I understood the concept. On the first stage I was not concentrated and swapped x:y, so I screwed up the whole result. During the breast cancer screening examination I did everything right. That went right into my head.

What really caused me problems is the naive Bayes classifier [17]. I have no idea why. I understand the formulas and know how to calculate everything, but still I keep stumbling across individual sections. I have to take a closer look at the topic and maybe watch some videos. Since I learned today with headphones and music, it could also be due to this. Anyway, I now know how the concept of a naive spam filter [18] works and I could program small in Python to understand that better.

All in all, the chapter was really interesting, but you can already see that here the level will be indicated and you have to exert yourself a little more. Music doesn't seem to help me to learn, but I really wanted to try it out (although I'm writing the article with music right now and it supports me here). One thing I really learned in this chapter today is that I really need to look more closely at the basic principles of statistics [19]. I know that this is very simple. I just need to understand a few places on the foundation so that I understand all the following formulas by myself. It's always like that when you work with numbers. Apart from that, I've done the good half and something that didn't suffer from my little mistakes are motivation and fun. That's a very good sign and slowly I'm getting warm with the different topics. I'm just so impatient and want to understand it all right away, but I have to accept that my brain needs enough time for that. I have set myself another personal challenge and will explain it in more detail at the end of the course.

This chapter was all about machine learning [20] and what can be added to the basis. There was a general introduction in which the whole machine learning voodoo was disenchanted. What journalists, marketing and advertising say about the topic is mostly unscientific bullshit. There's nothing magical, just numbers, dates and methods. The different types and for which applications accidents are used was quite interesting. Whether it's real estate or shopping. Once you have understood how exactly the method of the next neighbour [21] is used, for example, you wonder how understandable and logical it is. I got to know the MNIST database [22], which I would like to work on in the future.

Also where the difference between unsupervised [23] and supervised [24] classification methods lie is self-explanatory after a thorough study of the texts in this chapter. The simple introduction to simple linear regression [25] and logistic regression [26] is really well done and except for one small mistake I was able to solve all the tests without any problems. I noticed that the concept of diagrams works very well with my head because once I've seen formulas or examples in graphics, the information automatically flowed into my brain. That was exceedingly exciting to observe that with itself. There is not much more to say about this chapter, except that it was really fun. Very good structure and understandable examples.

The fifth part of the course dealt with neural networks [27] and how they are related to artificial intelligence. Basics such as deep learning [28], neurons, cell bodies and signals, for example, were covered. A brief look into the future takes a look, for example, at brain-computer interfaces [29] and how they could function. In the middle it is about the structure and how it is calculated. In my opinion, one position that could have been expanded further was the one with the levels. This can certainly be extended even further. Then we went from Conventional neural networks (cnns) [30] to generative adversarial networks [31]. The tests of the tasks were appealing and very interesting. Something that struck me in this chapter is that Finns have a very self-conscious and fine sense of humor. The last chapter dealt with the future, the societal implications of AI and the summary of the previous information.

As you can see above, I passed the course and also managed a challenge I had set myself. Normally you have 6 weeks, i.e. 42 days for the course. I was gonna finish the course in 24 days. If courses take too long for me, I quickly lose interest and start to neglect the course. I know my mistakes, so I always have to make it a little shorter to challenge myself. By my experience from my self-employment I can estimate quite well how long a project has to go and how I have to time that to be successful. Overall, I enjoyed the course very much. You had to read a lot of what I like. Videos (or even lectures) only make sure that I fall asleep quickly, no matter if the speaker can speak well. I need the information in book or paper form. My laptop and a room without any interference. Which is also important. Paper (one pen in red, one in black) to rethink important or complicated things. The design of Reaktor [32] is really pleasant and intuitive. A very big praise from me.

Conclusion

Which I also noticed, by the way. The Finnish government puts a lot of money into education [33] and there is a philosophy of eternal learning. Finland provides a very relaxed atmosphere and I am toying with the idea of doing the Java course [34]. Java [35] is a really horrible language and in the training as an IT specialist for application development I found it just terrible, but it would get me used to learning on my own and I could perhaps enroll in the right courses at Helsinki University [36] at some point. I was at the Technical University of Dortmund for two semesters and I didn't like it there. Everything is so rigid and boring, there is only information to be memorized to create the exams. After the examination the students are then on the job market and do not get a job, because nobody wants to pay a reasonable price for a computer scientist. The economy is wondering why nobody studies computer science anymore. Working here is boring, you live in an anti-digital atmosphere and the projects are ridiculous to backward. Even though Finland is a little too close to Russia for me, I would also take the hurdle of getting to know a new culture, language and people by making an offer or a study place.

Reference:

- [1] elementsofai.com/ (fetched: Fr, Jan 25, 2019 13:41)
- [2] helsinki.fi/en (fetched: Fr, Jan 25, 2019 13:42)
- [3] technologyreview.com/s/520431/driverless-cars-are-further-away-than-you-think/ (fetched: Fr, Jan 25, 2019 13:43)
- [4] web.archive.org/web/20151118212402/www-formal.stanford.edu/jmc/whatisai/whatisai.html (fetched: Fr, Jan 25, 2019 13:44)
- [5] cvpapers.com/index.html (fetched: Fr, Jan 25, 2019 13:45)
- [6] truesouthmarine.com/uncategorized/the-spacex-autonomous-spaceport-drone-ship/ (fetched: Fr, Jan 25, 2019 13:49)
- [7] harishnarayanan.org/writing/artistic-style-transfer/ (fetched: Fr, Jan 25, 2019 13:50)
- [8] makingscience.royalsociety.org/s/rs/people/fst00117605 (fetched: Fr, Jan 25, 2019 13:51)
- [9] psych.utoronto.ca/users/reingold/courses/ai/turing.html (fetched: Fr, Jan 25, 2019 13:52)
- [10] 7thavenueproject.com/post/117460013340/philosopher-john-searle-perception-mind-matter-conscious (fetched: Fr, Jan 25, 2019 13:53)
- [11] globetrotter.berkeley.edu/people/Searle/searle-con4.html (fetched: Fr, Jan 25, 2019 13:59)
- [12] xxx
- [13] web.archive.org/web/20131011125002/www-formal.stanford.edu/jmc/ (fetched: Fr, Jan 25, 2019 14:03)
- [14] mathworld.wolfram.com/MinimaxTheorem.html (fetched: Fr, Jan 25, 2019 14:04)
- [15] xxx
- [16] planetmath.org/BayesTheorem (fetched: Fr, Jan 25, 2019 14:06)
- [17] nlp.stanford.edu/IR-book/html/htmledition/naive-bayes-text-classification-1.html (fetched: Fr, Jan 25, 2019 14:07)
- [18] alexn.org/blog/2012/02/09/howto-build-naive-bayes-classifier.html (fetched: Fr, Jan 25, 2019 14:08)
- [19] psychstat3.missouristate.edu/Documents/IntroBook3/sbk.htm (fetched: Fr, Jan 25, 2019 14:09)
- [20] see.stanford.edu/Course/CS229 (fetched: Fr, Jan 25, 2019 14:10)
- [21] Peter Hall, et. al. Choice of Neighbor Order in nearest-Neighbor Classification, Australian National University, Souel National University and University of Cambridge, The Annals of Statistic, 2008, Vol 36, No. 5, 2135-2152, Institute of Mathematical Statistics 2008
- [22] yann.lecun.com/exdb/mnist/ (fetched: Fr, Jan 25, 2019 14:15)
- [23] mathworks.com/discovery/unsupervised-learning.html (fetched: Fr, Jan 25, 2019 14:16)
- [24] mathworks.com/discovery/supervised-learning.html (fetched: Fr, Jan 25, 2019 14:17)
- [25] people.duke.edu/~rnau/mathreg.htm (fetched: Fr, Jan 25, 2019 14:18)
- [26] omidrouhani.com/research/logisticregression/html/logisticregression.htm (fetched: Fr, Jan 25, 2019 14:19)
- [27] web.archive.org/web/20150726190848/deeplearning4j.org/neuralnet-overview.html (fetched: Fr, Jan 25, 2019 14:20)
- [28] Jürgen Schmidhuber et. al., Deep Learning in neural Networks: An Overview, Technical report IDSIA-03-14 / arXiv:1404.7828 v4 [cs.NE], 8 October 2014, University of Lugano, Switzerland
- [29] openbci.com/ (fetched: Fr, Jan 25, 2019 14:25)
- [30] ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/ (fetched: Fr, Jan 25, 2019 14:25)
- [31] Ian J. Godfellow et. al. Département d'informatique et de recherche opérationnelle Université de Montréal, QC H3C 3J7 2014
- [32] reaktor.com/ (fetched: Fr, Jan 25, 2019 14:30)
- [33] minedu.fi/en/frontpage (fetched: Fr, Jan 25, 2019 14:31)
- [34] mooc.fi/en/ (fetched: Fr, Jan 25, 2019 14:33)
- [35] xkcd.com/801/ (fetched: Fr, Jan 25, 2019 14:35)
- [36] courses.helsinki.fi/en/open-university/search/field_imp_teaching_language%253Avalue/en?search=&sorting=title_field%3Aasc&items_per_page=10&order=title_field&sort=asc (fetched: Fr, Jan 25, 2019 14:36)